

Online Combinatorial Optimization for Interconnected Refrigeration Systems: Linear Approximation and Submodularity

Insoon Yang*

Abstract

Commercial refrigeration systems consume 7% of the total commercial energy consumption in the United States. Improving their energy efficiency contributes to the sustainability of global energy systems and the supermarket business sector. This paper proposes a new control method that can save the energy consumption of multi-case supermarket refrigerators by explicitly taking into account their interconnected and switched system dynamics. Its novelty is a bilevel combinatorial optimization formulation to generate ON/OFF control actions for expansion valves and compressors. The inner optimization module keeps display case temperatures in a desirable range and the outer optimization module minimizes energy consumption. In addition to its energy-saving capability, the proposed controller significantly reduces the frequency of compressor switchings by employing a conservative compressor control strategy. However, solving this bilevel optimization problem associated with interconnected and switched systems is a computationally challenging task. To solve the problem in near real time, we propose two approximation algorithms that can solve both the inner and outer optimization problems at once. The first algorithm uses a linear approximation, and the second is based on the submodular structure of the optimization problem. Both are (polynomial-time) scalable algorithms and generate near-optimal solutions with performance guarantees. Our work complements existing optimization-based control methods (e.g., MPC) for supermarket refrigerators, as our algorithms can be adopted as a tool for solving combinatorial optimization problems arising in these methods.

Key words. Control systems, Refrigerators, Temperature control, Optimization, Integer linear programming, Greedy algorithms, Scalability

1 Introduction

Commercial refrigeration systems account for 7% of the total commercial energy consumption in the United States [29]. Therefore, there is a strong need for energy-efficient refrigeration systems, but research and development have focused on improving hardware rather than software, including control systems. Traditionally, hysteresis and set point-based controllers have been used to maintain the display case temperature in a desirable range without considering system dynamics and energy consumption. Over the past decade, however, more advanced control systems have been developed to save energy consumption using real-time sensor measurements and optimization algorithms (see Section 1.1). Advances in new technologies, such as the Internet of Things and cyber-physical systems, enhance the practicality of such an advanced control system with their sensing, communication, and computing capabilities [9].

*Ming Hsieh Department of Electrical Engineering, University of Southern California (insoonya@usc.edu). Supported in part by NSF under CRII:CPS (CNS1657100).




Figure 1: A supermarket refrigerator, which has 10 display cases. Evaporator i controls the temperature of display case i . Lines with arrows represent heat transfers between neighboring display cases or between a display case and ambient air.

Supermarkets are one of the most important commercial sectors in which energy-efficient refrigeration systems are needed. The primary reasons are twofold. First, supermarket refrigerators consume 56% of energy consumed by commercial refrigeration systems [17]. Second, supermarkets operate with very thin profit margins (on the order of 1%), and energy savings thus significantly help their business: the U.S. Environmental Protection Agency estimates that reducing energy costs by \$1 is equivalent to increasing sales by \$59 [1]. However, improving the energy efficiency of supermarket refrigerators is a challenging task because food products must be stored at proper temperatures. Failure to do so will increase food safety risks. The most popular refrigerators in supermarkets are multi-display case units. An example is illustrated in Fig. 1. Each display case has an evaporator controlled by an expansion valve, and a unit's suction pressure is controlled by a compressor rack, as shown in Fig. 2. In typical supermarket refrigerators, controllers turn ON and OFF expansion valves and compressors to keep display case temperatures in a specific range. Importantly, there are heat transfers between display cases due to the interconnection among them. Note that traditional hysteresis or set point-based controllers do not take into account such heat transfers and therefore perform in a suboptimal way.

This paper proposes a new control method that can improve the energy efficiency of multi-case supermarket refrigerators by explicitly taking into account the interconnected and switched dynamics of display case temperatures. The proposed controller receives sensor measurements and optimizes ON/OFF control actions for expansion valves and compressors in near real time. The novelty of this work is a bilevel combinatorial optimization formulation to generate such ON/OFF control signals in which (i) the inner combinatorial optimization module is responsible for maintaining display case temperatures in a desirable range, and (ii) the outer combinatorial optimization module minimizes energy consumption. The primary advantage of the proposed approach is its energy savings. Because the controller explicitly takes into account the system dynamics and heat transfers, it effectively uses state measurements and optimizes control actions to save energy while guaranteeing desired temperature profiles. In our case studies, the proposed control method saves 7.5–8% of energy compared to a traditional approach. The secondary benefit of the proposed method is to reduce the frequency of compressor switchings. It is known that frequent switchings of compressors accelerate their mechanical wear. We propose a conservative compressor control approach that reduces fluctuations in suction pressure and thus decreases the compressor switching frequency. In our case studies using a benchmark refrigeration system model, the proposed method reduces the switching frequency by 54–71.6%.

The proposed control method, however, presents a theoretical and algorithmic challenge because a bilevel combinatorial optimization associated with a dynamical system must be solved in near real time. To overcome this challenge, we suggest two approximation algorithms that can solve both of the inner and outer optimization problems at once. The first algorithm uses the linear approximation method developed in our previous work [32]. The approximate problem is a linear




Figure 2: Schematic diagram of a supermarket refrigerator.

binary program, which can be solved by an efficient and scalable single-pass algorithm. In addition, it simulates the dynamical system model only once to generate control actions at each time point. We also show that the approximate solution obtained by this method has a provable suboptimality bound. The second algorithm is based on the *submodular* structure in the optimization problem. The inner optimization’s objective function is submodular because opening an expansion valve when a smaller set of valves are opened gives a greater marginal benefit than opening it when a larger set of valves are already opened. We prove this intuitive submodularity property. Therefore, a greedy algorithm can be adopted to obtain a $(1 - \frac{1}{e})$ -optimal solution [18]. In our case studies, the actual performance of the proposed controller using these two algorithms is 98.9–99.5% of the optimal controller.

1.1 Related Work

Several optimization-based control methods for commercial refrigerators have been developed over the past decade. One of the most popular methods is model predictive control (MPC) although it is computationally challenging to apply standard MPC due to the switched dynamics of refrigeration systems. It is shown that the mixed logical dynamical framework is useful to solve small-size problems with a piecewise affine approximation of a system model [4, 12]. However, the practicality of this method is questionable due to the high dimensionality of practical problems for supermarket refrigerators, except for limited cases. To overcome this limitation, [24] carefully selects and parametrizes optimization variables to formulate the problem as nonlinear MPC instead of hybrid MPC. Nonetheless this approach is computationally expensive because a nonlinear program with many variables must be solved in each MPC iteration. An alternative approach using hierarchical MPC is proposed in [28]. This method separates time scales into two: in every nonlinear MPC iteration, low-level temperature controllers were employed, and the high-level optimization task is to determine optimal parameters for these controllers. However, this approach still presents the combinatorial growth of the search space. More recently, a sequential convex programming-based method is shown to be computationally efficient in several case studies [10]. It iteratively solves an optimization problem using convex programming, replacing the nonconvex cost function with a convex approximation. In several numerical experiments, this heuristic method generates high-quality control signals although it gives no theoretical performance guarantee. We believe that our work is complementary to the aforementioned methods. One of our main contributions is to develop two efficient and scalable algorithms for resolving the computational challenge in discrete optimization problems associated with supermarket refrigeration systems. These algorithms can be

adopted as a tool for solving combinatorial optimization problems in the aforementioned methods. We propose one of the most efficient control architectures that use the algorithms.

With advances in modern power systems, an important emerging application is using supermarket refrigeration systems for thermal storage [19, 25, 30, 16]. In this application, it is often necessary to regulate the total power consumption of hundreds of refrigerators to provide reliable demand response services to the power grid. Our work contributes to such demand response applications by providing scalable online optimization algorithms with performance guarantees that are useful for solving large-scale problems. The utility of the proposed algorithms in demand response is demonstrated in [31] and a case study is presented in Section 5.

1.2 Outline

The remainder of this paper is organized as follows. In Section 2, we describe an interconnected hybrid system model of supermarket refrigerators and provide a simulation result when a traditional set point- and PI-based controller is employed. We then propose the proposed control method based on bilevel online combinatorial optimization in Section 3. In Section 4, we provide two efficient algorithms to solve the combinatorial optimization problem near real time and examine their scalability. In Section 5, we compare the performance of the proposed controllers with that of the traditional controller and demonstrate its utility in automated demand response.

2 Switched and Interconnected Dynamics of Supermarket Refrigeration Systems

We consider a supermarket refrigerator in which multiple display cases are interconnected with one another. For example, Fig. 1 shows a refrigerator that has 10 display cases. The temperature of each display case is controlled by an evaporator unit, where the refrigerant evaporates absorbing heat from the display case. Let evaporator i be in charge of display case i for $i = 1, \dots, n$, where n is the number of display cases in all the refrigerators. Several dynamic models of supermarket refrigeration systems have been proposed [23, 13, 14, 21, 22, 26] (see also the references therein). Among those, we use the benchmark model of a typical supermarket refrigeration system proposed in [13] and widely used in [24, 28, 34, 30, 16]. This model is useful for simulating display case temperatures and evaluating the performances of several controllers.

2.1 Display Cases and Evaporators

Display cases store food products and keep them refrigerated. This refrigeration is due to the heat transfer between the food product and the cold air in the display cases. Let $T_{\text{food},i}$ and $T_{\text{air},i}$ denote the temperatures of the food product and the air in display case i . The heat transfer $Q_{\text{food} \rightarrow \text{air},i}$ between the food product and the air in display case i can then be modeled as

$$\begin{aligned} m_{\text{food},i} c_{\text{food},i} \dot{T}_{\text{food},i} &= -Q_{\text{food} \rightarrow \text{air},i} \\ &= -k_{\text{food-air}}(T_{\text{food},i} - T_{\text{air},i}), \end{aligned} \quad (2.1)$$

where $m_{\text{food},i}$ is the mass of the food product, $c_{\text{food},i}$ is the heat capacity of the food product and $k_{\text{food-air}}$ is the heat transfer coefficient between the food product and the air.

The display case air temperature is affected by the heat transfers from the food product ($Q_{\text{food} \rightarrow \text{air},i}$), the ambient air ($Q_{\text{amb} \rightarrow \text{air},i}$), the evaporator ($-Q_{\text{air} \rightarrow \text{evap},i}$) and the neighboring display case air ($\sum_{j=1}^n Q_{j \rightarrow i}$). The refrigerant flow into an evaporator is controlled by its expansion

valve. Let u_i be the valve control variable for evaporator i such that

$$u_i(t) := \begin{cases} 0 & \text{if expansion valve } i \text{ is closed at } t \\ 1 & \text{otherwise.} \end{cases}$$

Expansion valve i controls the refrigerant injection into evaporator i and decreases the pressure of the refrigerant if it is open, as shown in Fig. 2. Then, the dynamics of the display case air temperature can be modeled as the following switched interconnected system:

$$\begin{aligned} m_{\text{air},i} c_{\text{air},i} \dot{T}_{\text{air},i} &= Q_{\text{food} \rightarrow \text{air},i} + Q_{\text{amb} \rightarrow \text{air},i} - Q_{\text{air} \rightarrow \text{wall},i} + \sum_{j=1}^n Q_{j \rightarrow i} \\ &= k_{\text{food-air}}(T_{\text{food},i} - T_{\text{air},i}) + k_{\text{amb-air}}(T_{\text{amb}} - T_{\text{air},i}) \\ &\quad - k_{\text{air-evap}}(T_{\text{air},i} - T_{\text{evap}} u_i) + \sum_{j=1}^n k_{i,j}(T_{\text{air},j} - T_{\text{air},i}), \end{aligned} \tag{2.2}$$

where T_{amb} is the ambient air temperature, T_{evap} is the refrigerant's evaporation temperature, $k_{\text{amb-air}}$ is the heat transfer coefficient between the ambient air and the display case air and $k_{i,j}$ is the heat transfer coefficient between display case i 's air and display case j 's air. Note that $k_{i,j} = 0$ if display cases i and j are not neighbors. For a more detailed model, one can separately consider the dynamics of the evaporator wall temperature [24].¹ However, the proposed model is a good approximation because the heat transfer coefficient between the evaporator wall and the refrigerant is five to ten times higher than other heat transfer coefficients [13].

The mass flow out of the evaporator can be computed as

$$f_i := \frac{1}{\Delta t} m_{\text{ref},i},$$

where the refrigerant mass in the evaporator is controlled by the valve switching

$$m_{\text{ref},i} = \begin{cases} m_{\text{ref}}^{\text{max}} & \text{if } u_i = 1 \\ 0 & \text{if } u_i = 0. \end{cases}$$

Depending on the specification of refrigerators, it takes a nontrivial amount of time to fill up the evaporator by refrigerant. In this case, the dynamics of the refrigerant mass in the evaporator can be explicitly taken into account [24]. Alternatively, one can introduce a delay-time constant, τ , and let $m_{\text{ref},i}(t) = m_{\text{ref}}^{\text{max}} u_i(t - \tau)$ to model the effect of the time to fill up the evaporator.

2.2 Suction Manifold and Compressor Rack

As shown in Fig. 2, the evaporated refrigerant with low pressure from the outlet of the evaporator is compressed by the electric motors in the compressor bank. Each refrigerator could have multiple compressors and each compressor is switched ON or OFF. For example, all the compressors are turned ON when maximal compression is needed. The compressor bank is conventionally controlled by a PI controller to maintain the suction pressure within a bandwidth.

¹Alternatively, one can introduce a delay parameter, τ , and replace $T_{\text{evap}} u_i(t)$ with $T_{\text{evap}} u_i(t - \tau)$ to explicitly take into account the effect of the evaporator wall temperature.

The suction manifold pressure P_{suc} evolves with the following dynamics:

$$\dot{P}_{\text{suc}} = \frac{1}{V_{\text{suc}} r_{\text{suc}}} \left(\sum_{i=1}^n f_i - \rho_{\text{suc}} \sum_{i=1}^{n_c} F_{c,i} \right), \quad (2.3)$$

where V_{suc} is the volume of the suction manifold, ρ_{suc} is the density of the refrigerant in the suction manifold, and $r_{\text{suc}} := d\rho_{\text{suc}}/dP_{\text{suc}}$. The variable $F_{c,i}$ denotes the volume flow out of the suction manifold controlled by compressor i . Let $u_{c,i}$ be the control variable for compressor i , where $u_{c,i} = 0$ represents that compressor i is OFF and $u_{c,i} = 1$ represents that compressor i is ON. The volume flow $F_{c,i}$ is then given by

$$F_{c,i} = k_c u_{c,i} := \frac{\eta V_{\text{comp}}}{n} u_{c,i},$$

where η is the volumetric efficiency of each compressor, and V_{comp} denotes the compressor volume.

The total power consumption by the compressor rack is given by

$$p = \rho_{\text{suc}} (h_{oc} - h_{ic}) \sum_{i=1}^{n_c} F_{c,i},$$

where h_{ic} and h_{oc} are the enthalpies of the refrigerant flowing into and out of the compressor, respectively. The compressed refrigerant flows to the condenser and is liquefied by generating heat, as shown in Fig. 2. The liquefied refrigerant flows to the expansion valve, and as a result, the refrigeration circuit is closed.

2.3 Traditional Set-Point/PI-Based Control

A widely used control method consists of (i) a set-point based control of expansion valves, and (ii) a PI control of compressors [13, 24]. Specifically, the following ON/OFF control law is traditionally used for expansion valve i :

$$u_i(t) := \begin{cases} 1 & \text{if } T_{\text{air},i}(t) > T_i^{\text{max}} \\ 0 & \text{if } T_{\text{air},i}(t) < T_i^{\text{min}} \\ u_i(t^-) & \text{otherwise,} \end{cases}$$

where $[T_i^{\text{min}}, T_i^{\text{max}}]$ is the desirable temperature range for display case i . To control the suction pressure, compressors are traditionally operated by a PI controller. This controller tracks the error $e(t)$ that measures the deviation of the suction pressure from the reference \bar{P}_{suc} over the dead band DB , i.e.,

$$e(t) := \begin{cases} P_{\text{suc}}(t) - \bar{P}_{\text{suc}} & \text{if } |e(t)| > DB \\ 0 & \text{otherwise.} \end{cases}$$

Then, the number of ON compressors is determined by a thresholding rule depending on the following output of the PI controller

$$u_{PI}(t) = K_P e(t) + \frac{1}{K_I} \int e(t) dt;$$

the greater value the output generates, the more compressors the controller turns on. More details about the thresholding rule can be found in [13, 24].

In our case studies, R134a is chosen as the refrigerant. Its relevant thermodynamic properties are contained in [13]. For convenience, we summarize the properties as follows: $T_{\text{evap}} = -4.3544 P_{\text{suc}}^2 +$

$29.2240P_{\text{suc}} - 51.2005$, $\Delta h = (0.0217P_{\text{suc}}^2 - 0.1704P_{\text{suc}} + 2.2988) \times 10^5$, $\rho_{\text{suc}} = 4.6073P_{\text{suc}} + 0.3798$, $r_{\text{suc}} = -0.0329P_{\text{suc}}^3 + 0.2161P_{\text{suc}}^2 - 0.4742P_{\text{suc}} + 5.4817$, $\rho_{\text{suc}}(h_{oc} - h_{ic}) = (0.0265P_{\text{suc}}^3 - 0.4346P_{\text{suc}}^2 + 2.4923P_{\text{suc}} + 1.2189) \times 10^5$. These formulas were obtained by fitting polynomials with experimental data. The additional parameters used in the simulations are summarized in Table 1.

Table 1: Parameters used in simulations

Parameter	Value	Parameter	Value
$\bar{m}_{\text{food},i}$	200 kg	$c_{\text{food},i}$	1000 J/(kg·K)
$m_{\text{wall},i}$	260 kg	$c_{\text{wall},i}$	385 J/(kg·K)
$m_{\text{air},i}$	50 kg	$c_{\text{air},i}$	1000 J/(kg·K)
$k_{\text{food-air}}$	300 J/(s·K)	$k_{\text{air-evap}}$	225 J/(s·K)
$k_{i,j}$	500 J/(s·K)	$k_{\text{amb-evap}}$	275 J/(s·K)
$m_{\text{ref}}^{\text{max}}$	1 kg	V_{suc}	10 m ³
η	0.81	V_{comp}	0.2 m ³ /s
n	10	T_{amb}	20°C
T_i^{min}	0°C	T_i^{max}	5°C
K_P	0.1	K_I	-0.8
P_{suc}	1.4 bar	DB	0.3 bar

We perturbed the mass of food products in each display case by $\pm 20\%$ from the nominal value $\bar{m}_{\text{food},i}$. Despite this heterogeneity, the set point-based controller almost identically turns ON and OFF all the expansion valves and therefore all the display case temperatures have almost the same trajectory as shown in Fig. 3 (a). This synchronization is due to the decentralized nature of the set point-based controller: the control decision for expansion valve i depends only on its local temperature. Intuitively, this decentralized controller is suboptimal because it does not actively take into account the heat transfer between neighboring display cases. This inefficiency of the traditional control approach motivates us to develop a new optimization-based control method that explicitly considers the interdependency of display case temperature dynamics.

Another disadvantage resulting from the synchronization of expansion valves is the significant fluctuation of suction pressure. Since the PI controller integrates the deviation of suction pressure from its reference, the output $u_{PI}(t)$ presents large and frequent variations. As a result, the number of ON compressors frequently varies as shown in Fig. 3 (b). A frequent switching of compressors is a serious problem because it accelerates the mechanical degradation of the compressors. Our strategy to discourage frequent compressor switchings is twofold: (i) our conservative compressor control method tries to maintain $P_{\text{suc}}(t) = \bar{P}_{\text{suc}}$, not fully utilizing the pressure bandwidth $\pm DB$, and (ii) our online optimization-based controller indirectly desynchronizes the ON/OFF operation of expansion valves. The details about the two control methods are presented in the following sections. Unlike the traditional control approach, our proposed method is suitable for regulating the total power consumption in real time. This feature is ideal for modern power system (so-called ‘smart grid’) applications, allowing supermarket units to follow a real-time regulation signal for reducing peak demand or supporting a spinning reserve. Such applications of our control method to power systems are studied in [31] and one of which is studied in Section 5.4.




Figure 3: (a) The food temperatures in display cases 1–5, operated by the PI controller over 8 hours. The temperatures in all display cases are almost identical. (b) The number of ON compressors operated by the PI controller. Note that the profile presents frequent fluctuations.

3 Control via Online Combinatorial Optimization

3.1 Conservative Compressor Control

We control the compressor rack to (approximately) maintain the suction pressure as the reference \bar{P}_{suc} , i.e.,

$$P_{\text{suc}}(t) \approx \bar{P}_{\text{suc}} \quad \forall t.$$

In other words, to make $\dot{P}_{\text{suc}} \equiv 0$ in (2.3), we set $u_c := (u_{c,1}, \dots, u_{c,n_c})$ such that the refrigerant outflow from the suction manifold is equal to the inflow:

$$\rho_{\text{suc}} \sum_{i=1}^{n_c} k_c u_{c,i} \approx \sum_{i=1}^n f_i. \quad (3.1)$$

In practice, we may not be able to exactly satisfy this equality because each $u_{c,i}$ is either 0 or 1. However, we assume that the compressor control action u_c can be chosen to make the difference between the outflow and the inflow negligible. This compressor control rule is suboptimal: it induces a conservative operation of the compressor rack that does not fully utilize the pressure bandwidth. However, this conservative control approach has a practical advantage: it does not create significant compressor switchings. Therefore, it can potentially decelerate the mechanical wear of compressors. Under this compressor control rule, the total power consumption can be computed as

$$\begin{aligned} p &= (h_{oc} - h_{ic}) \sum_{i=1}^n f_i \\ &= \frac{(h_{oc} - h_{ic}) m_{\text{ref}}^{\max}}{\Delta t} \sum_{i=1}^n u_i. \end{aligned} \quad (3.2)$$

3.2 Bilevel Optimization Formulation

We consider a receding-horizon online optimization approach to generate control signals for expansion valves and compressors. Let $\{t_0, t_1, \dots, t_k, t_{k+1}, \dots\}$ be the time steps at which the control

action is optimized. For the sake of simplicity, we describe a one-step look-ahead optimization method; however, this approach can be easily extended to multiple-step look-ahead optimization (see Remark 2).

3.2.1 Inner problem for temperature management

At time t_k , we control the expansion valves to minimize the following quadratic deviation from the upper-bound T_i^{\max} , $i = 1, \dots, n$:

$$J(\alpha) = \sum_{i=1}^n \int_{t_k}^{t_{k+1}} (T_{\text{air},i} - T_i^{\max})_+^2 dt,$$

where $(a)_+^2 = a^2 \cdot \mathbf{1}_{\{a \geq 0\}}$, assuming T_{air} is evolving with (2.1) and (2.2). Specifically, the expansion valve action at t_k is generated as a solution to the following combinatorial optimization problem:

$$\min_{\alpha \in \{0,1\}^n} J(\alpha) \quad (3.3a)$$

$$\text{s.t. } \dot{x} = Ax + Bu + C, \quad x(t_k) = \mathbf{x}_{\text{meas}} \quad (3.3b)$$

$$u(t) = \alpha, \quad t \in (t_k, t_{k+1}] \quad (3.3c)$$

$$\|\alpha\|_0 = \sum_{i=1}^n \alpha_i \leq K. \quad (3.3d)$$

Here, $x := (T_{\text{food}}, T_{\text{air}})$ and (3.3b) gives a linear system representation of the dynamics (2.1) and (2.2). Note that \mathbf{x}_{meas} represents $(T_{\text{food}}, T_{\text{air}})$ measured at $t = t_k$.² As specified in (3.3c), the control action over $(t_k, t_{k+1}]$ is fixed as the solution α . The last constraint (3.3d) essentially limits the power consumed by the refrigeration system as $K(h_{oc} - h_{ic})m_{\text{ref}}^{\max}/\Delta t$ due to (3.2). Therefore, the choice of K is important to save energy: as K decreases, the power consumption lessens.

3.2.2 Outer problem for energy efficiency

To generate an energy-saving control action, we minimize the number K of open expansion valves while guaranteeing that the quadratic deviation $J(\alpha)$ from the upper-bound T_i^{\max} , $i = 1, \dots, n$ is bounded by the threshold Δ . More precisely, we consider the following outer optimization problem:

$$\min\{K \in \{0, \dots, n\} \mid J(\alpha^{\text{opt}}(K)) \leq \Delta\}, \quad (3.4)$$

where $\alpha^{\text{opt}}(K)$ is a solution to the expansion valve optimization problem (3.3). Let K^{opt} be a solution to this problem. Then, $\alpha^{\text{opt}}(K^{\text{opt}})$ is the expansion valve control action that saves energy the most while limiting the violation of the food temperature upper-bound T_i^{\max} , $i = 1, \dots, n$. This outer optimization problem can be easily solved by searching K from 0 in an increasing order. Once we find \hat{K} such that $J(\alpha^{\text{opt}}(\hat{K})) \leq \Delta$, we terminate the search and obtain the solution as $K^{\text{opt}} := \hat{K}$. In the following section, we will show that this procedure can be integrated into approximation algorithms for the inner optimization problem.

Then, as specified in (3.3c), the controller chooses $u^{\text{opt}}(t) := \alpha^{\text{opt}}(K^{\text{opt}})$ for $t \in (t_k, t_{k+1}]$. Furthermore, it determines the compressor control signal u_c^{opt} such that $\sum_{i=1}^{n_c} u_{c,i}^{\text{opt}} \approx m_{\text{ref}}^{\max}/(\rho_{\text{suc}} k_c \Delta t)$ using (3.1). If $P_{\text{suc}}(t) < \bar{P}_{\text{suc}}$, the controller rounds $m_{\text{ref}}^{\max}/(\rho_{\text{suc}} k_c \Delta t)$ to the next smaller integer and then determines the number of ON compressors as the integer. If $P_{\text{suc}}(t) \geq \bar{P}_{\text{suc}}$, the controller rounds $m_{\text{ref}}^{\max}/(\rho_{\text{suc}} k_c \Delta t)$ to the nearest integer greater than or equal to it. The information flow in this control system is illustrated in Fig. 4.

²If T_{food} is not directly measured, an observer needs to be employed to estimate the state. Then, the control system uses the estimate $T_{\text{food}}^{\text{est}}$ instead of its actual measurement as shown in Fig. 4.




Figure 4: The proposed control system with the outer and inner optimization modules.

Remark 1. Our objective function $J(\alpha)$ only takes into account the violation of temperature upper-bounds. This choice is motivated by the fact that the food temperature in each display case increases as we close more expansion valves, which is summarized in Proposition 1. In other words, as we reduce the number K of open valves in the outer optimization problem, the possibility of violating temperature upper-bounds increases, while it is less likely to violate temperature lower-bounds. This monotonicity property of food temperatures justifies our focus on temperature upper-bounds.

Proposition 1. Let $T_{food,j}^\alpha$ and $T_{air,j}^\alpha$ denote the food and air temperatures in display case j when the control action α is applied. Then, for any $\alpha, \beta \in \mathbb{R}^n$ such that

$$\alpha_i \leq \beta_i, \quad i = 1, \dots, n,$$

we have

$$T_{food,j}^\alpha \geq T_{food,j}^\beta \text{ and } T_{air,j}^\alpha \geq T_{air,j}^\beta, \quad j = 1, \dots, n.$$

Proof. The proof is contained in Appendix A. □

4 Approximation Algorithms

We present two approximation methods for the inner optimization problem. One is based on linear approximation, and another utilizes submodularity. These will give approximate solutions with guaranteed suboptimality bounds. We further show that, by simply modifying these approximation algorithms for the inner optimization problem, we can obtain a near-optimal solution to the outer optimization problem.

4.1 Linear Approximation-Based Optimization

We first consider a linear approximation-based approach to the inner combinatorial optimization problem (3.3). It is convenient to work with the following value function:

$$V(\alpha) = J(0) - J(\alpha). \tag{4.1}$$

The value $V(\alpha)$ represents the reduction in the quadratic deviation from from the upper-bound T_i^{\max} , $i = 1, \dots, n$, when expansion valve j is chosen to be open only for j such that $\alpha_j = 1$. Note that this value function is normalized such that $V(0) = 0$. The Taylor expansion of V at 0 gives $V(\alpha) = DV(0)^\top \alpha + O(\alpha^2)$ assuming the derivative DV is well-defined. This motivates us to

consider the following first-order approximation of the expansion valve optimization problem (3.3):

$$\begin{aligned} \max_{\alpha \in \{0,1\}^n} \quad & DV(0)^\top \alpha \\ & \|\alpha\|_0 \leq K. \end{aligned} \quad (4.2)$$

The i th entry $[DV(0)]_i$ of the derivative represents the marginal benefit of opening expansion valve i . Therefore, the approximate problem (4.2) can be interpreted as maximizing the marginal benefit of valve operation while guaranteeing that the number of open valves is less than or equal to K . A detailed method to define and compute the derivative can be found in [32]. Computing the derivative should also take into account the dependency of the state x on the binary decision variable α . For example, an adjoint-based approach can be used to handle this dependency [11].

The first advantage of the proposed approximation approach is that it gives an approximate solution with a provable suboptimality bound. The bound is *a posteriori*, which does not require the globally optimal solution α^{opt} but the solution α^* of (4.2).

Theorem 1 ([32]). *Let α^* be a solution to the approximate problem (4.2). If $DV(0)^\top \alpha^* \neq 0$, then the following suboptimality bound holds:*

$$\rho V(\alpha^{opt}) \leq V(\alpha^*),$$

where

$$\rho = \frac{V(\alpha^*)}{DV(0)^\top \alpha^*} \leq 1.$$

If $DV(0)^\top \alpha^* = 0$, then $V(\alpha^{opt}) = V(0) = 0$, i.e., 0 is an optimal solution.

Its proof is contained in Appendix B. This theorem suggests that the approximate solution's performance is greater than $(\rho \times 100)\%$ of the globally optimal solution's performance.

The second advantage of the proposed method is that it yields an efficient algorithm to solve the approximate problem (4.2). Specifically, we design a very simple algorithm based on the ordering of the entries of $DV(0)$. Let $\mathbf{d}(\cdot)$ denote the map from $\{1, \dots, n\}$ to $\{1, \dots, n\}$ such that

$$[DV(0)]_{\mathbf{d}(i)} \geq [DV(0)]_{\mathbf{d}(j)} \quad (4.3)$$

for any $i, j \in \{1, \dots, n\}$ such that $i \leq j$. Such a map can be constructed using a sorting algorithm with $O(n \log n)$ complexity. Such a map may not be unique. We let $\alpha_{\mathbf{d}(i)} = 1$ for $i = 1, \dots, K$ if $[DV(0)]_{\mathbf{d}(i)} > 0$. A more detailed algorithm to solve this problem is presented in Algorithm 1. Note that it is a single-pass algorithm, i.e., does not require multiple iterations. Furthermore, Lines 9–12 can be parallelized.

Remark 2. *The proposed linear approximation method is applicable to multi-period optimization problems, in which the objective is given by $J(\alpha) := \sum_{k=1}^{N_{period}} J_k(\alpha^k)$ and the control variable is time-varying, i.e., $\alpha = (\alpha^1, \dots, \alpha^{N_{period}}) \in \mathbb{R}^{n \times N_{period}}$. In such a case, we compute the derivative DV_k of $V_k(\alpha^k) := J_k(0) - J_k(\alpha^k)$ for each k . The objective function can be approximated as $\sum_{k=1}^{N_{period}} DV_k(0)^\top \alpha^k$, which is still linear in α . Therefore, we can use the proposed algorithm.*

4.2 Submodular Optimization

The second approach gives another approximate solution of the expansion valve optimization problem (3.3) with a suboptimality bound. This solution is generally different from the solution obtained

Algorithm 1: Algorithm for the approximate problem (4.2)

```

1 Initialization:
2  $\alpha \leftarrow 0$ ;
3 Construction of  $\mathbf{d}$ :
4 Compute  $DV(0)$ ;
5 Sort the entries of  $DV(0)$  in descending order;
6 Construct  $\mathbf{d} : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$  satisfying (4.3);
7 Solution of (4.2):
8 while  $[DV(0)]_{\mathbf{d}(i)} > 0$  and  $i \leq K$  do
9    $\alpha_{\mathbf{d}(i)} \leftarrow 1$ ;
10   $i \leftarrow i + 1$ ;
11 end

```

by the first approach. Let $\Omega := \{1, \dots, n\}$ be the set of expansion valves to be controlled. We define a set function, $\mathcal{V} : 2^\Omega \rightarrow \mathbb{R}$, as

$$\mathcal{V}(X) = V(\mathbb{I}(X)),$$

where the value function V is defined as (4.4) and $\mathbb{I}(X) := (\mathbb{I}_1(X), \dots, \mathbb{I}_n(X)) \in \{0, 1\}^n$ is the indicator vector of the set X such that $\mathbb{I}_i(X) := 0$ if $i \notin X$ and $\mathbb{I}_i(X) := 1$ if $i \in X$. In other words, \mathcal{V} is a set function representation of V . The expansion valve optimization problem (3.3) is equivalent to selecting the set $X \subseteq \Omega$ such that $|X| \leq K$ to maximize the value function $\mathcal{V}(X)$, i.e.,

$$\begin{aligned} \max_{X \in 2^\Omega} \quad & \mathcal{V}(X) \\ \text{s.t.} \quad & |X| \leq K. \end{aligned} \tag{4.4}$$

We observe that the value function \mathcal{V} has a useful structure, which is called the *submodularity*. It represents a *diminishing return* property such that opening an expansion valve when a smaller set of valves is opened gives a greater marginal benefit than opening it when a larger set of valves is already opened.

Theorem 2. *The set function $\mathcal{V} : 2^\Omega \rightarrow \mathbb{R}$ is submodular, i.e., for any $X \subseteq Y \subseteq \Omega$ and any $\mathbf{a} \in \Omega \setminus Y$,*

$$\mathcal{V}(X \cup \{\mathbf{a}\}) - \mathcal{V}(X) \geq \mathcal{V}(Y \cup \{\mathbf{a}\}) - \mathcal{V}(Y).$$

Furthermore, it is monotone, i.e., for any $X \subseteq Y \subseteq \Omega$

$$\mathcal{V}(X) \leq \mathcal{V}(Y).$$

Proof. See Appendix C. □

The submodularity of \mathcal{V} guarantees that Algorithm 2, which is a greedy algorithm, provides an $(1 - \frac{1}{e})$ -optimal solution. In other words, the approximate solution's performance is greater than $(1 - \frac{1}{e}) \approx 63\%$ of the oracle's performance. In our case study, the actual submodularity is 98.9%, which is significantly greater than this theoretical bound.

Theorem 3 ([18]). *Algorithm 2 is a $(1 - \frac{1}{e})$ -approximation algorithm. In other words, if we let X^* be the solution obtained by this greedy algorithm, then the following suboptimality bound holds:*

$$\left(1 - \frac{1}{e}\right) \mathcal{V}(X^{opt}) \leq \mathcal{V}(X^*),$$

Algorithm 2: Greedy algorithm for (4.4)

```

1 Initialization:
2  $X \leftarrow \emptyset$ ;
3 Greedy algorithm:
4 while  $i \leq K$  do
5    $\mathbf{a}^* \in \arg \max_{\mathbf{a} \in \Omega \setminus X} \mathcal{V}(X \cup \{\mathbf{a}\})$ ;
6    $X \leftarrow X \cup \{\mathbf{a}^*\}$ ;
7    $i \leftarrow i + 1$ ;
8 end

```

where X^{opt} is an optimal solution to (4.4).

Lines 5–9 of Algorithm 2 makes a locally optimal choice at each iteration. Therefore, it significantly reduces the search space, i.e., it does not search over all possible combinations of open expansion valves. When the expansion valve optimization problem (3.3) is extended to multi-stage optimization, a greedy algorithm achieves the same suboptimality bound using adaptive (or string) submodularity and monotonicity of the value function [8, 2, 15].

4.3 Modified Algorithms for the Outer Optimization Problem

We now modify the two approximation algorithms for the inner problem (3.3) to solve the full bilevel optimization problem. In both Algorithms 1 and 2, the expansion valve chosen to be open at iteration i (line 10 of Algorithm 1 and line 7 of Algorithm 2) is *independent of the selections at later iterations*. This independency plays an essential role in incorporating the outer optimization problem into the algorithms. To be more precise, we compare the cases of $K = l$ and $K = l + 1$. Let α^l and α^{l+1} be the solutions in the two cases obtained by Algorithm 1. Since the expansion valve selected to be open at iteration $l + 1$ do not affect the choices at earlier iterations, we have $\alpha_{\mathbf{d}(i)}^l = \alpha_{\mathbf{d}(i)}^{l+1}$ for $i = 1, \dots, l$. Therefore, we do not have to re-solve the entire inner optimization problem for $K = l + 1$ if we already have the solution for $K = l$; it suffices to run one more iteration for $i = l + 1$ to obtain $\alpha_{\mathbf{d}(l+1)}^{l+1}$. This observation allows us to simply modify lines 8–12 in Algorithm 1 as Algorithm 3. As we can see in line 2, we select expansion valves to be open until when the

Algorithm 3: Modified version of Algorithm 1 for the outer optimization problem (3.4)

```

1 while  $[DV(0)]_{\mathbf{d}(i)} > 0$  and  $J(\alpha) > \Delta$  do
2    $\alpha_{\mathbf{d}(i)} \leftarrow 1$ ;
3    $i \leftarrow i + 1$ ;
4 end

```

temperature upper-bound violation $J(\alpha)$ is less than or equal to the threshold Δ . Similarly, we modify lines 5–9 of Algorithm 2 as Algorithm 4 to solve the outer problem.

4.4 Scalability

We now compare the complexity of Algorithms 3 and 4. Algorithm 3, which is based on a linear approximation, is *single pass* in the sense that, after computing the derivative and ordering its entries only once, we can obtain the solution. Calculating the derivative requires $O(n^2 N_T)$, where

Algorithm 4: Modified version of Algorithm 2 for the outer optimization problem (3.4)

```

1 while  $J(\mathbb{I}(X)) > \Delta$  do
2    $\mathbf{a}^* \in \arg \max_{\mathbf{a} \in \Omega \setminus X} \mathcal{V}(X \cup \{\mathbf{a}\});$ 
3    $X \leftarrow X \cup \{\mathbf{a}^*\};$ 
4    $i \leftarrow i + 1;$ 
5 end

```

N_T is the number of time points in the time interval $[t_k, t_{k+1}]$ used to integrate the dynamical system [32], if a first-order scheme is employed. Therefore, the total complexity including the sorting step is $O(n^2 N_T) + O(n \log n)$. On the other hand, Algorithm 4, which is a greedy algorithm, chooses a locally optimal solution at each stage. In other words, this iterative greedy choice approach requires one to find an entry that maximizes the increment in the current payoff at every stage. Its complexity is $O(n^3 N_T)$. Therefore, Algorithm 3 is computationally more efficient as the number n of display cases grows. Note, however, that Algorithm 4 is also scalable because its complexity is cubic in n and linear in N_T .

5 Case Studies

In this section, we examine the performance of the proposed online optimization-based controllers. For fair comparisons with the traditional controller, we use the parameter data reported in Section 2.3 with a refrigerator unit that has 10 display cases. Fig. 5 and 6 illustrates the simulation results of the proposed controllers.

5.1 Energy Efficiency

As opposed to the synchronized food temperature profiles controlled by the traditional method (see Fig. 3 (a)), the proposed controllers induce alternating patterns of the temperatures as shown in Fig. 5. Such patterns result from the explicit consideration of heat transfers between neighboring display cases in the optimization module through the constraint (3.3b), which represents the interconnected temperature dynamics. Using the spatial heat transfers, the proposed controllers do not turn ON or OFF all the expansion valves at the same time. Instead, they predict the temperature evolution for a short period and selects the valves to turn ON that are effective to minimize the deviation from the desirable temperature range during the period. As a result, the ON/OFF operation of expansion valves is desynchronized, unlike in the case of the traditional controller. This desynchronization maintains the temperatures near the upper-bound T^{\max} reducing temperature fluctuations. Therefore, it intuitively improves energy efficiency. As summarized in Table 2, the proposed controllers save 7.5–8% of energy.

Table 2: Energy savings by the proposed controllers

	PI	linear	submodular
average kW	11.24	10.34	10.40
energy saving	–	8.0%	7.5%
suboptimality	90.1%	99.5%	98.9%

Note that the outer optimization module minimizes the total energy consumption while the inner optimization module is responsible for maintaining the temperature profiles in a desirable




Figure 5: The food temperatures (in 5 display cases out of 10) controlled by (a) the linear approximation-based algorithm (Algorithm 3), and (b) the submodular optimization algorithm (Algorithm 4).




Figure 6: The number of ON compressors controlled by (a) the linear approximation-based algorithm (Algorithm 3), and (b) the submodular optimization algorithm (Algorithm 4).

range. When the bilevel combinatorial problem is exactly solved for all time, the average power consumption is 10.29kW. Therefore, the two proposed controllers' performances are 99.5% and 98.9% of the optimal controller although their theoretical suboptimality bounds are 39% and 63%.

5.2 Reduced Compressor Switching

Another advantage of the proposed controllers is the considerable reduction on the number of compressor switching instances. By desynchronizing the switching instances of expansion valves in the inner optimization module, the proposed controllers significantly reduce the variation of suction pressure. Our conservative compressor control approach presented in Section 3.1 also helps to minimize the deviation of the suction pressure from its reference. As a result, the controllers significantly reduce the fluctuations on the number of ON compressors as shown in Fig. 6. First, the maximum number of ON compressors is decreased from six to two. This reduction suggests that a mechanically more compact compressor or a smaller number of compressors in the rack may be enough if the proposed controllers are adopted. Second, the proposed controllers reduce the number of compressor switching instances by 54.0–71.6% as summarized in Table 3. These

infrequent compressor operation strategies are beneficial to decelerate the mechanical degradation of compressors.

Table 3: Compressor switching reductions by the proposed controllers

	PI	linear	submodular
# of switchings	324	92	149
reduction	–	71.6%	54.0%

5.3 Comparisons of the Two Proposed Controllers

Fig. 5 illustrates that the submodular optimization-based controller maintains the temperatures in a narrower range than the linear approximation-based controller. This feature is owing to the fact that the greedy algorithm used in the submodular optimization-based method avoids violating the temperature upper bound in a locally optimal way. However, to keep the temperatures in a narrower range this approach requires a faster adjustment of suction pressure than the linear approximation-based method. As a result, the proposed compressor controller performs a more frequent switching of compressors when the submodular optimization-based method is adopted (see Fig. 6 and Table 3). Furthermore, this frequent compressor switching induces an inefficient use of the compressor rack and therefore turns ON more compressors on average (in time). Therefore, the submodular optimization-based controller consumes slightly more energy than the linear approximation-based controller as reported in Table 2.

5.4 Automated Demand Response under Real-Time Pricing

Real-time pricing of electricity refers to passing wholesale prices through to end users. At least in theory, it is shown to improve electricity market efficiency among other benefits [5].³ However, consumer should bear the risk of receiving high energy bills if consumers do not appropriately react to fluctuating wholesale prices. Such a risk transfer to end-users under real-time pricing can be reduced by automated demand response (ADR) [20] and can also be limited by contracts for ADR [33]. In this subsection, we demonstrate the utility of our method as a control tool for refrigeration ADR under real-time pricing. In particular, we consider the scenario of energy arbitrage: supermarket refrigeration systems automatically consume less energy when the real-time price is high and consume more when it is low. Note that real-time prices are often difficult to predict and hence ADR must be performed in an online fashion by appropriately reacting to fluctuating prices. The online optimization feature of our controllers allows them to adjust energy consumption in response to real-time prices (by changing the number K of ON expansion valves in real time). Fig. 7 (a) shows the real-time wholesale electricity price at the Austin node in the Electric Reliability Council of Texas (ERCOT) on July 3, 2013 [7]. We use a simple thresholding law for choosing K : if the electricity price is greater than \$0.1/kWh, the controller allows up to 70% of expansion valves to turn ON; otherwise, it operates as usual. As summarized in Table 4, the proposed controllers save 14.3–15.0% of energy cost compared to a standard PI controller. In addition, the temperature deviations from T^{\max} are less than 0.5°C as shown in Fig. 7 (b), because right after the reduction in energy consumption the controller encourages enough cooling to recover the desired temperature levels.⁴ Further applications of the proposed algorithms to ADR

³However, real-time prices fail to capture the economic value of responsive loads in general [27].

⁴Such post-cooling is mostly feasible due to the mean-reverting behavior of electricity prices which discourages sustained price peaks [6]. We can also perform pre-cooling if prices are predictable or their distributional information is available. Such pre-cooling will increase the economic benefit of the proposed method.




Figure 7: (a) The real-time price data; The food temperatures (in five display cases out of ten) controlled by (b) the linear approximation-based algorithm (Algorithm 3), and (c) the submodular optimization algorithm (Algorithm 4).

aggregating a large number of supermarket refrigerator units can be found in [31].

Table 4: Operation costs per refrigerator under real-time pricing from 10am to 6pm

	PI	linear	submodular
cost	\$5.67	\$4.82	\$4.86
cost saving	—	15.0%	14.3%

6 Conclusions

The proposed controller explicitly takes into account the switched and interconnected dynamics, and is therefore suitable for multi-case supermarket refrigeration systems. However, it has to solve a bilevel combinatorial optimization problem associated with switched interconnected systems in near real time, which is a challenging task. To overcome this difficulty, we proposed two polynomial-time approximation algorithms that are based on the structural properties of this optimization problem. These algorithms can also be adopted as a tool for solving combinatorial optimization problems arising in existing MPC-based methods. We demonstrated the performance of the proposed controllers through case studies using a benchmark refrigeration system model and found that (i) they improve energy efficiency by 7.5–8%, (ii) they reduce the number of compressor switchings by 54–71.6%, and (iii) they save 14.3–15% of operation cost under a demand response scenario. In addition to conventional usages, the scalability of the proposed algorithms can contribute to an emerging methodology of controlling a large number of refrigerator units through cloud computing.

A Proof of Proposition 1

Proof. We use the linear system representation (3.3b) of the food and air temperature dynamics (equations (2.1) and (2.2)). We first notice that

$$A_{i,j} \geq 0 \quad \forall i \neq j,$$

where $A_{i,j}$ represents the (i,j) th entry of the matrix A . Furthermore, $k_{\text{air-evap}}T_{\text{evap}} \leq 0$ due to the non-positive evaporator temperature. Hence, we have

$$B_{i,j} \leq 0 \quad \forall i, j.$$

Using Proposition III.2 in [3], we conclude that the system (3.3b) is *input-monotone* such that for any $\alpha, \beta \in \mathbb{R}^n$ with $\alpha_i \leq \beta_i$, $i = 1, \dots, n$,

$$x_i^\alpha \geq x_i^\beta, \quad i = 1, \dots, n,$$

where x^α denotes the solution of the system (3.3b) when its input is chosen as α . \square

B Proof of Theorem 1

Proof. In (3.3b), we notice that

$$x(t) = e^{A(t-t_k)}\mathbf{x}_{\text{meas}} + \int_{t_k}^t e^{A(t-s)}B\alpha ds,$$

which implies that $x(t)$ is linear in α . Therefore, V is concave with respect to α in a continuously relaxed space, \mathbb{R}^n . Then, the result follows from Theorem 2 in [32]. \square

C Proof of Theorem 2

Proof. Let $T_{\text{food},i}^X$ denote the temperature of the food product in display case i given that the expansion valves in X are open. Due to the linearity of the system dynamics (2.1) and (2.2), $T_{\text{food},i}^X$ is modular, i.e.,

$$T_{\text{food},i}^X = \sum_{\mathbf{a} \in X} T_{\text{food},i}^{\{\mathbf{a}\}}.$$

Therefore, for any $X \subseteq Y \subseteq \Omega$ and any $\mathbf{a} \in \Omega \setminus Y$

$$T_{\text{food},i}^{X \cup \{\mathbf{a}\}} - T_{\text{food},i}^X = T_{\text{food},i}^{Y \cup \{\mathbf{a}\}} - T_{\text{food},i}^Y.$$

Furthermore, Proposition 1 yields the following monotonicity result: for any $X \subseteq Y \subseteq \Omega$

$$T_{\text{food},i}^X \geq T_{\text{food},i}^Y,$$

i.e., as we open more expansion valves, the food temperature decreases. Lastly, the concavity of $\mathcal{V}(X) = \mathcal{V}(\emptyset) - \sum_{i=1}^n \int_{t_k}^{t_{k+1}} (T_{\text{food},i}^X - T_i^{\max})_+^2 dt$ in $T_{\text{food},i}^X$ implies that $X \subseteq Y \subseteq \Omega$ and any $\mathbf{a} \in \Omega \setminus Y$

$$\mathcal{V}(X \cup \{\mathbf{a}\}) - \mathcal{V}(X) \geq \mathcal{V}(Y \cup \{\mathbf{a}\}) - \mathcal{V}(Y).$$

Therefore, \mathcal{V} is submodular. Its monotonicity follows from Proposition 1. \square

References

- [1] *ENERGY STAR Building Upgrade Manual Chapter 11: Supermarkets and Grocery Stores*, 2008.
- [2] S. Alaei and A. Malekian. Maximizing sequence-submodular functions and its application to online advertising. *arXiv:1009.4153 [cs.DM]*, 2010.
- [3] D. Angeli and E. D. Sontag. Monotone control systems. *IEEE Transactions on Automatic Control*, 48(10):1684–1698, 2003.
- [4] A. Bemporad and M. Morari. Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3):407–427, 1999.
- [5] S. Borenstein and S. P. Holland. On the efficiency of competitive electricity markets with time-invariant retail prices. *RAND Journal of Economics*, 36(3):469–493, 2005.
- [6] A. Cartea and M. G. Figueroa. Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance*, 12(4):313–335, 2005.
- [7] Electric Reliability Council of Texas. Real-time prices reports.
- [8] D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.
- [9] M. Graziano and M. Pritoni. Gloudfridge: A cloud-based control system for commercial refrigeration systems. Technical report, ACEEE Summer Study on Energy Efficiency in Buildings, 2014.
- [10] T. G. Hovgaard, S. Boyd, L. F. S. Larsen, and J. B. Jørgensen. Nonconvex model predictive control for commercial refrigeration. *International Journal of Control*, 86(8):1349–1366, 2013.
- [11] P. Kokotović and J. Heller. Direct and adjoint sensitivity equations for parameter optimization. *IEEE Transactions on Automatic Control*, 12(5):609–610, 1967.
- [12] L. F. S. Larsen, T. Geyer, and M. Morari. Hybrid model predictive control in supermarket refrigeration systems. In *Proceedings of 16th IFAC World Congress*, 2005.
- [13] L. F. S. Larsen, R. Izadi-Zamanabadi, and R. Wisniewski. Supermarket refrigeration system - benchmark for hybrid system control. In *Proceedings of 2007 European Control Conference*, 2007.
- [14] B. Li and A. G. Alleyne. A dynamic model of a vapor compression cycle with shut-down and start-up operations. *International Journal of Refrigeration*, 33(3):538–552, 2010.
- [15] Y. Liu, E. K. P. Chong, A. Pezeshki, and B. Moran. Bounds for approximate dynamic programming based on string optimization and curvature. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, 2014.
- [16] T. Minko, R. Wisniewski, J. D. Bendtsen, and R. Izadi-Zamanabadi. Cost efficient optimization based supervisory controller for supermarket subsystems with heat recovery. In *Proceedings of 2015 European Control Conference*, 2015.

- [17] Navigant Consulting, Inc. Energy savings potential and R&D opportunities for commercial refrigeration. Technical report, U.S. Department of Energy, 2009.
- [18] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximising submodular set functions – I. *Mathematical Programming*, 265–294, 1978.
- [19] R. Pedersen, J. Schwensen, S. Sivabalan, C. Corazzol, S. E. Shafiei, K. Vinther, and J. Stoustrup. Direct control implementation of a refrigeration system in smart grid. In *Proceedings of 2013 American Control Conference*, 2013.
- [20] M. A. Piette, S. Kiliccote, and G. Ghatikar. Design and implementation of an open, interoperable automated demand response infrastructure. Technical report, Lawrence Berkeley National Laboratory, 2008.
- [21] B. P. Rasmussen. Dynamic modeling for vapor compression systems–part i: Literature review. *HVAC&R Research*, 18(5):934–955, 2012.
- [22] B. P. Rasmussen. Dynamic modeling for vapor compression systems–part ii: Simulation tutorial. *HVAC&R Research*, 18(5):956–973, 2012.
- [23] B. P. Rasmussen and A. G. Alleyne. Dynamic modeling and advanced control of air conditioning and refrigeration systems. Technical report, Air Conditioning and Refrigeration Center, University of Illinois at Urbana-Champaign, 2006.
- [24] D. Sarabia, F. Capraro, L. F. S. Larsen, and C. de Prada. Hybrid NMPC of supermarket display cases. *Control Engineering Practice*, 17:428–441, 2009.
- [25] S. E. Shafiei, R. Izadi-Zamanabadi, H. Rasmussen, and J. Stoustrup. A decentralized control method for direct smart grid control of refrigeration systems. In *Proceedings of the 52nd IEEE Conference on Decision and Control*, 2013.
- [26] S. E. Shafiei, H. Rasmussen, and J. Stoustrup. Modeling supermarket refrigeration systems for demand-side management. *Energies*, 6(2):900–920, 2013.
- [27] R. Sioshansi and W. Short. Evaluating the impacts of real time pricing on the usage of wind power generation. *IEEE Transaction on Power Systems*, 24(2):516–524, 2009.
- [28] C. Sonntag, A. Devanathan, and S. Engell. Hybrid NMPC of a supermarket refrigeration system using sequential optimization. In *Proceedings of 17th IFAC World Congress*, 2008.
- [29] U.S. Department of Energy. Commercial real estate energy alliance. 2012.
- [30] K. Vinther, H. Rasmussen, J. Stoustrup, and A. G. Alleyne. Learning-based precool algorithms for exploiting foodstuff as thermal energy reserve. *IEEE Transactions on Control Systems Technology*, 23(2):557–569, 2015.
- [31] I. Yang. *Risk Management and Combinatorial Optimization for Large-Scale Demand Response and Renewable Energy Integration*. PhD thesis, University of California, Berkeley, 2015.
- [32] I. Yang, S. A. Burden, R. Rajagopal, S. S. Sastry, and C. J. Tomlin. Approximation algorithms for optimization of combinatorial dynamical systems. *IEEE Transactions on Automatic Control*, 2016.

- [33] I. Yang, D. S. Callaway, and C. J. Tomlin. Indirect load control for electricity market risk management via risk-limiting dynamic contracts. In *Proceedings of 2015 American Control Conference*, pages 3025–3031, 2015.
- [34] Z. Yang, K. B. Rasmussen, A. T. Kieu, and R. Izadi-Zamanabadi. Fault detection and isolation for a supermarket refrigeration system-part one: Kalman-filter-based methods. In *Proceedings of 18th IFAC World Congress*, 2011.